Chapter 24

ON THE USE OF DATA ASSIMILATION IN BIOGEOCHEMICAL MODELLING

Andreas Oschlies
School of Ocean and Earth Science, Southampton Oceanography Centre, Southampton, UK

Abstract A main objective of applying data assimilation methods to marine ecosystem models is the optimisation of often poorly known model parameters or even of the model's functional form. Recent efforts in this direction are reviewed. Results obtained so far indicate that presently available data sets can constrain not more 10 to 15 different ecological parameters. This raises questions about the use of more complex models. On the other hand, none of the optimised models yields a satisfactory fit to the observations, suggesting that present ecosystem models are overly simplistic. Implications of these apparently contradictory findings are discussed and a data assimilative strategy for future improvement of marine ecosystem models is suggested.

Keywords: Marine ecosystem models, data assimilation, parameter optimisation.

1. Introduction

Interest in prognostic models of marine biogeochemical cycles arises to a large extent from our need to better understand, quantify, and eventually predict the ocean's role in the global carbon cycle. This includes cycles of related elements, such as nitrogen, phosphorus or iron, that can act as limiting nutrients for phytoplankton growth. Other aspects addressed by biogeochemical and ecological modelling include the prediction of harmful algal blooms [*Schofield et al.*, 1999], and a quantitative understanding of oceanic food webs up to fish [*Loukos et al.*, 2003], birds, and humans, as well as the possible impact of marine sulfur emissions on the formation of cloud condensation nuclei [*Gabric et al.*, 2004]. In this chapter, I will focus on the carbon issue.

Carbon fluxes in the ocean are often described in terms of solubility pump and biological pump. The abiotic solubility pump is caused by increasing solubility of CO_2 (as of other gases) with decreasing temperature. For present climate conditions, deep water forms at high latitudes and average ocean temperatures are colder than average sea surface temperatures. The solubility pump then ensures that the volume averaged carbon concentration is larger than the surface averaged one, yielding a vertical CO_2 (or, more precisely, dissolved inorganic carbon, DIC) gradient with higher concentrations at depth than at the surface (Figure 1).

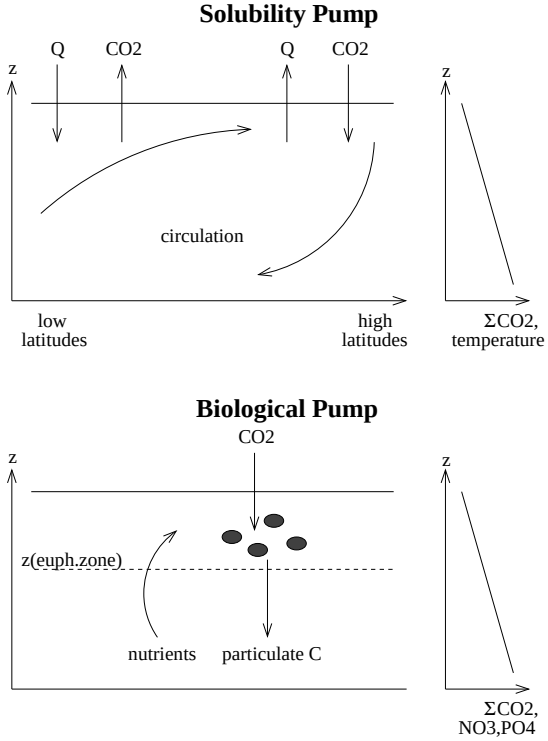

Figure 1. Schematic representation of the solubility pump (top) and the biological pump (bottom), both acting to maintain the vertical gradient in total dissolved inorganic carbon (ΣCO_2) in the ocean. Q is the surface heat flux, with oceanic heat uptake corresponding to outgassing and cooling to CO_2 uptake, z(euph.zone) is the depth of the euphotic zone which describes the surface region with light levels sufficient to allow for photosynthesis (typically about 100 m).

The term "pump" reflects that carbon is transported against the mean vertical gradient. Closer analysis reveals that of the presently observed

vertical DIC gradient, only about a quarter can be explained by the solubility pump [*Sarmiento et al.,* 1995] with the remaining three quarters being attributed to the "biological" carbon pump [*Volk and Hoffert,* 1985]. The driving agent of the biological pump is photosynthesis that generates organic carbon and thereby reduces DIC concentrations, and accordingly the partial pressure of CO_2, in the surface ocean. Respiration of organic carbon by metabolic processes in bacteria, higher trophic levels, and in the photosynthetically active phytoplankton itself reverses this process. As a result of mixing and advection along the vertical light gradient and because of the formation of biogenic particles that sink through the water instead of moving with it, respiration occurs generally deeper in the water column than photosynthesis. This decoupling of photosynthesis and respiration generates vertical gradients of DIC. To make things more complicated, some organisms form calcium carbonate "hard parts" which, on formation, sinking, and dissolution also affect the carbonate chemistry of sea water and result in an alkalinity pump. Because the formation of calcium carbonate in surface water increases surface pCO_2, this constitutes a counter pump in terms of CO_2 which partly compensates the pCO_2 effect of the organic carbon pump. A robust mechanistic understanding of the formation and biotically aided dissolution of calcium carbonate shells is not yet available, and many models so far assume that a fixed fraction of all biogenic particulate carbon sinking out of the light-lit euphotic zone (roughly the upper 100 m) is associated with calcium carbonate formation.

A close interaction of biology and physics arises not only from the interplay of physical and biological transport mechanisms on the vertical DIC gradient, but also from the fact that phytoplankton growth requires the presence of both light and nutrients, which usually have opposite vertical gradients. Accordingly, light and nutrient levels experienced by a phytoplankton cell are very sensitive to physical transport processes that may upwell or entrain deeper and nutrient-rich waters, or may mix or advect phytoplankton cells down into the dark ocean interior. This physical control on biological production has to be taken into account when attempting to simulate the marine carbon cycle. A standard strategy is to couple marine ecosystem models into circulation models. Validation of such coupled models is not straightforward. For example, the strong sensitivity of the marine biota to physical transport processes makes it difficult to separately evaluate the individual model components. For many applications one can at least safely neglect the biological impact on the physics via changes in the absorption profile of solar radiation [*Oschlies,* 2004]. While this allows to evaluate the physical model component individually, the reverse is not true for the

impact of the physics on the marine biogeochemistry. Here, a potential mapping of errors of the physical model onto the predicted ecosystem fields makes the separate evaluation of the ecosystem model component difficult. This is not necessarily a disadvantage: Because of the marine biology's strong sensitivity and fast response time of the order of days to changes in the light or nutrient supply, coupling ecosystem and circulation models may actually help to identify deficiencies of physical transport processes, particularly in the upper ocean [e.g., *Oschlies*, 1999].

The following section will give a brief overview over the field of biogeochemical models and presently used marine ecosystem models. Section 3 discusses some aspects of observations that are relevant for data assimilation, and section 4 addresses the potential use of combining data and biogeochemical models. Data assimilation methods are presented in section 5, and this article ends with a discussion of some achievements and perspectives of data assimilation in the field of biogeochemical modelling.

2. Biogeochemical modelling

Compared to numerical modelling of the ocean circulation, the field of biogeochemical modelling is much less mature. In particular, there is no known equivalent to the Navier-Stokes equations. In principle, these describe the motion of sea water exactly, but an exact solution does not (yet?) exist. The rules are thus clear for physical models, and different numerical models basically attempt to find different approximations to the unknown exact solution.

The situation is different in the field of biogeochemical modelling. Although there are some reliable, albeit mainly empirical, laws that describe transformations among various inorganic compounds dissolved in sea water, things become relatively shaky once life, and thereby transformations among organic and inorganic chemical compounds, comes into play. In practice, biogeochemical models are generally composed of an inorganic chemistry component and an ecosystem component, of which the latter is the by far more complex, expensive, and problematic part. In the following I will focus on this ecosystem model component and often use the term ecosystem model as synonym for the entire biogeochemical model.

Marine ecosystem models usually attempt to describe life's action on marine biogeochemical tracers by partitioning the marine ecosystem into a handful of boxes, often called compartments, such as phytoplankton (plants), zooplankton (animals), or detritus (dead organic matter). Sometimes, a further distinction is made between particular and dis-

solved dead organic matter which does not sink but moves passively with the water. Besides the different transport properties, the distinction among dissolved and particulate organic matter is also useful in terms of elemental ratios: while the elemental composition of particulate organic matter is, on average, found to be close to the Redfield ratio [*Redfield*, 1934; *Redfield et al.*, 1963], dissolved organic matter often contains several times more carbon than the Redfield ratio would predict [*Williams*, 1995; *Kähler and Koeve*, 2001].

Using mass conservation as underlying concept, the compartments simulate stocks of atoms of the relevant element, and fluxes such as primary production, grazing, or mortality all describe the transfer of atoms among the different compartments. Often only a single element (usually one associated with a potentially limiting nutrient, e.g., nitrogen for nitrate, phosphorus for phophate) is modelled explictly, and its concentration in each of the compartments becomes a prognostic variable. Concentrations and fluxes of other elements (in particular, carbon) are usually diagnosed via the Redfield ratio. While this seems to be consistent with the analysis of average inorganic remineralisation products [e.g., *Anderson and Sarmiento*, 1995], more detailed investigations reveal local and temporal systematic deviations [*Körtzinger et al.*, 2001; *Sterner and Elser*, 2002; *Klausmeier et al.*, 2004]. A few recent models have therefore begun to explicitly resolve the cycling of different elements [*Moore et al.*, 2002].

For each marine ecosystem model, the particular choice of its compartments and of the parameterisation of fluxes between the compartments contains subjective elements, which may for example be influenced by operational measurement protocols, historical paradigms, or taxonomic nomenclature. Such an approach is, of course, valid and probably necessary in a field in which a strong theoretical framework is not yet available (and in which key species may not even be discovered yet). Progress will be made by trying to construct models that can explain the observations and at the same time tell a plausible story, and by more or less steadily changing the story as new observations add new information. In this process it is important to keep in mind that the underlying rules that make up a particular ecosystem model are generally assumed rather than demonstrated and hence are subject to change.

After these cautionary remarks about the theoretical foundations of marine ecoystem models, it is time to point out that these very models may greatly help to improve our understanding of marine ecosystems by allowing us to test the assumed hypothetical rules against observations in a quantitative way. In the following I will try to present my subjective view of how this can be achieved in practice.

2.1 Ecosystem model types

Today, a large variety of marine ecosystem models exist, probably similar in number to the number of researchers in the field. Although strict categories do not exist, present models roughly fall into three main groups (Figure 2):

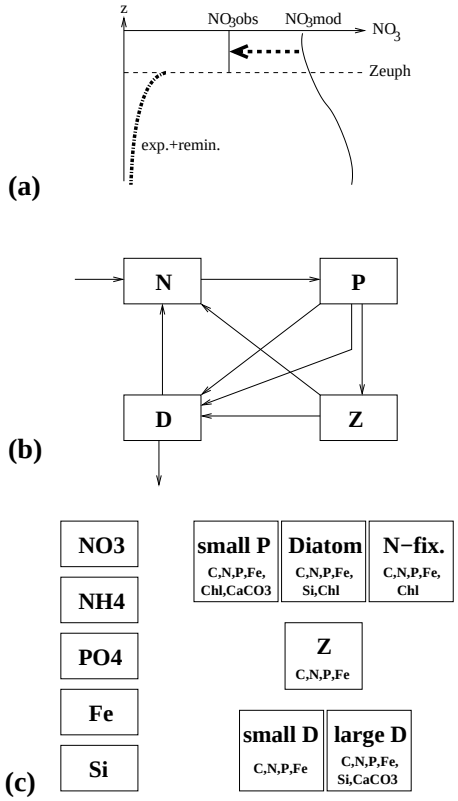

Figure 2. Schematic representation of various ecosystem model concepts: (a) Restoring of nutrients to observed or zero surface concentrations and immediate export and remineralisation according to a prescribed vertical remineralisation profile (e.g., *Bacastow and Maier-Reimer* [1990]). (b) Nutrient-Phytoplankton-Zooplankton-Detritus (NPZD) model (e.g., *Oschlies and Garçon* [1999]). (c) Multi-element multi-functional group model (after *Moore et al.* [2002]). Each biological compartment is composed of sub-compartments for each of the prognostic elements and chemical compounds indicated. For clarity of the illustration, the $O(100)$ fluxes among the various (sub)-compartments are not shown.

- Nutrient-restoring models. These models do not explicitly resolve ecosystem components other than a (usually) single nutrient. Bi-

ological production is simulated by restoring to either zero or observed nutrient concentrations in the light-lit surface layer, and instant sinking and remineralisation are accounted for by a prescribed remineralisation profile. Examples are the models of *Bacastow and Maier-Reimer* [1990], *Najjar et al.* [1992] and the models used during phases 1 and 2 of the Ocean Carbon Model Intercomparison Project (OCMIP, [*Orr*, 1999]). Depending on whether or not dissolved organic matter is explicitly resolved, these biogeochemical/ecosystem models typically have 2 to 4 parameters. They are most widely used in models that address time scales much longer than a year, and applications to seasonal or shorter time scales will be problematic because of the absence of any particulate organic-matter storage pools.

- NPZD-type models. Although NPZD stands for Nutrient, Phytoplankton, Zooplankton, and Detritus, such models may contain a few more prognostic variables like bacteria or dissolved organic matter. Most of these models are descendants of a configuration proposed by *Fasham et al.* [1990] and they explicitly simulate the cycling of either nitrogen or phosphorus. They have been applied to general ocean circulation models ranging from coarse resolution [*Sarmiento et al.*, 1993; *Fasham et al.*, 1993; *Chai et al.*, 1996; *McCreary et al.*, 1996; *Six and Maier-Reimer*, 1996] to eddy-permitting [*Oschlies and Garçon*, 1998, 1999] and eddy-resolving resolution [*Oschlies*, 2002]. Typically, these ecosystem models have 10 to 30 parameters.

- Functional-group type models. Going beyond the NPZD-type structure, these recently emerging models attempt to resolve different species or groups of phytoplankton and zooplankton. According to their special ecological function (e.g., nitrogen fixation, calcification) these are often called functional groups. These different groups require (and allow) to explicitly resolve the cycling of different biogeochemical elements. Examples are the models described by *Moore et al.* [2002] and *Aumont et al.* [2004], as well as the European Regional Seas Ecosystem Model (ERSEM, *Ebenhöh et al.* [1997]) and the evolving Dynamic Green Ocean Model [*Le Quéré et al.*, Ecosystem dynamics based on plankton functional types for global ocean biogeochemistry models, submitted to *Global Change*]. They typically have far more than hundred parameters.

3. Observations

A trivial statement is that the ocean is severly undersampled and that we need more data to better understand what is going on out there. We also have data of very different quality. There is a large number of data which are difficult to interpret in terms of ecological variables or processes resolved by dynamical models. Examples include wet zooplankton weight, satellite ocean colour data (which contain information on water-leaving radiance at a few wavelengths, but not immediately on surface chlorophyll or even primary production), or uptake of carbon isotopes into particulate matter (which is related but not identical to primary production, e.g., *Dring and Jewson* [1982]). Particular care has to be taken when different measurement methods that attempt –and often claim– to measure the same quantity (e.g., chlorophyll concentration, primary production) in fact measure different things . In contrast to more straightforward measurements of concentrations of standing stocks of organic or inorganic matter, direct observations of processes or rates (e.g., growth, grazing, sinking, exudation, mortality) are usually difficult to carry out without perturbing the system under investigation and, accordingly, are very limited in number and often have large random and systematic errors.

Available observations are also often biased towards the spring and summer season, with generally very few ship-based winter or autumn observations, particularly in mid and high latitudes. The same sampling bias holds for measurements of physical variables, but may be more critical for ecological properties for which the amplitude of the local seasonal cycle can be as large as the global range of the respective annual mean property.

Valuable observational information can also be taken from laboratory studies. Investigations using cultured species may, for example, help to reveal physiological information on the impact of environmental conditions like nutrient concentrations, light intensity, turbulence levels, or temperature on growth rates. A caveat to be kept in mind is that those species that have been cultured so far are not necessarily representative of the open-ocean plankton community. Considering that the number of generations separating our domestic plants and animals from their wild ancestors is reached by phytoplankton in only a few years, culture species may also be affected by selection and mutation.

It appears that information on the loss side (e.g., grazing, mortality) is more difficult to obtain than on the production side (growth). There is a (perhaps related?) tendency of marine ecosystem modellers to increase model complexity preferentially on the production side rather than on

the loss side. The net impact of marine biology on biogeochemical cycles is, however, controlled by the balance of production and loss processes. Because marine phytoplankton seems to invest relatively more into defence structures (mineral cell walls, spines, chains and colonies) than do land plants, which seem to compete more for fastest growth, one might even argue that marine ecosystems are more sensitive to loss processes than are terrestrial ecosystems [*Smetacek,* 2001].

4. Motivation for data assimilation

In a situation in which our understanding of marine ecosystem dynamics is still relatively poor and in which observations and data types are distributed unevenly, data assimilation may be seen as promising tool to interpolate in time and space as well as among different data types. Dynamical, albeit hypothetical, rules, e.g. in form of model equations, help to go beyond purely statistical interpolation and to link the observational information according to these rules. As such models have various poorly known parameters and functional relationships, data assimilation can at the same time provide a vehicle to estimate these parameters and parameterisations. This is conceptually different from state estimation that attempts to find a model state that agrees best with the observations and possibly a previous model forecast.

State estimation is used frequently in the field of meteorology to initialise new forecast simulations. For marine biogeochemistry, this aspect is generally less relevant, although it has already been applied for operational planning of research cruises [*Popova et al.,* 2002]. Forecast times are typically limited to a few weeks. The dissipative character of the dynamics that we believe to hold for marine ecosystems and that we use in our models [*Popova et al.,* 1997] and the strong seasonal and intraseasonal forcing in form of light, temperature, and mixing regimes lead to a quick memory loss of the initial conditions in typically much less than a year.

With respect to improving longer term forecasts, e.g. for climate prediction purposes, it seems to be more promising to rely on parameter estimation (and "parameterisation estimation") to improve our quantitative understanding of marine ecosystem dynamics. Data assimilation then provides a tool to test various hypothetical model dynamics against the available observations in an organised and quantitative way. The following sections attempt to give some overview over recent activities in this area.

5. Data assimilation methods used in marine biogeochemical modelling

5.1 Sequential methods

Sequential assimilation methods are constructed to accumulate information gathered from both observations and model predictions in time with the aim to generate an optimal state estimate. This approach is widely used in operational systems for which fast and robust delivery of information is a crucial aspect, ranging from instrument-guided aircraft landing, and related less peaceful applications, to atmospheric and oceanic weather forecasting. Most methods are approximations to, or descendants of, the Kalman Filter [*Kalman*, 1960]. Its basic principle is that of optimal interpolation between an observation and its simulated counterpart.

To illustrate its concept, we start by considering a model with a single state variable, x, for which observational counterparts are available. For Gaussian errors of the observation, x_{obs}, and model forecast, x_f, the best linear unbiased estimate (BLUE) of the true state vector is then given by:

$$x_a = \frac{\frac{x_f}{\sigma_f} + \frac{x_{obs}}{\sigma_{obs}}}{\frac{1}{\sigma_f} + \frac{1}{\sigma_{obs}}} \tag{24.1}$$

where σ_f and σ_{obs} are expected rms errors of model forecast and observation, respectively. The expected rms error of the analysed state, x_a, is

$$\sigma_a = \frac{1}{\frac{1}{\sigma_f} + \frac{1}{\sigma_{obs}}} \quad . \tag{24.2}$$

The fact that σ_a is smaller than both σ_f and σ_{obs} is consistent with x_a containing more information about the true state vector than any of the model forecast and observation alone. The analysed state, x_a, can then be used as initial condition to integrate the model until the next observation becomes available and the above process is repeated. Note, however, that the model forecast error will, in general, not be constant in time. For example, it may be large during particulate phases of the annual cycle (e.g., during the spring bloom), or it may be smaller when a lot of observations have been assimilated in the recent past. Computing the evolution of the model forecast error is the main contribution of the Kalman Filter. For a linear model, this is achieved in the following way:

$$\sigma^2_{n+1} = \mathbf{A}\sigma^2_n\mathbf{A}^t + \mathbf{e}^2_n \tag{24.3}$$

where $\mathbf{A}$ is the matrix that computes the evolution of the state vector $\mathbf{x}$, composed of the individual prognostic model variables, from time step

n to $n+1$ via $\mathbf{x}_{n+1} = \mathbf{A}\mathbf{x}_n$, $\mathbf{A}^t$ is its transpose, and $\mathbf{e}$ is some intrinsic uncertainty of the model that at each time step increases the forecast error. This equation is written in vector form to point out the main computational burden of the Kalman Filter: for an N-dimensional state vector $\mathbf{x}$, σ^2 is a $N \times N$ matrix. This means that the computation of σ^2_{n+1} is as expensive as stepping forward the state vector $2N$ times. In many applications, updating the covariance matrix will dominate the computational effort and often render it impractical.

Note, that, in addition to computational constraints, the above approach will provide optimal solutions only for linear systems and for Gaussian error distributions. Both conditions are not generally met for marine biogeochemical systems, and a number of adaptations to the original Kalman Filter approach have been developed. To cope with nonlinear systems, the so-called Extended Kalman Filter (EKF) steps forward the error covariance equation by the tangent linear operator (Jacobian) of the full model operator [*Evensen,* 1992]. While this works well for weakly nonlinear systems, the Ensemble Kalman Filter (EnKF) approach [*Evensen,* 2003] can also be applied to strongly nonlinear systems. It uses Monte-Carlo generated state vectors to initialise an ensemble of model forecast runs from which the error distribution of the model forecast is estimated. More recent developments include the Singular Evolutive Extended Kalman (SEEK) filter that has been applied to assimilate ocean colour data by *Carmillet et al.* [2001]. A sequential method that is not directly related to the Kalman Filter is the Sequential Importance Resampling (SIR) filter [*Bertino et al.,* 2003].

A common feature of all sequential methods is the generation of a model trajectory that is "only" piecewise self-consistent. Whenever observations become available, merging the respective model forecast and observation into a new analysed state generates unsteady "jumps" in the state vector trajectory. Special care has to be taken if one wants, for example, to ensure mass conservation across these jumps. Similarly, any analysis of output from an assimilation experiment has to account for fluxes or perturbations associated with the assimilation steps.

5.2 Variational methods

While sequential assimilation methods attempt to estimate a "best" state vector at each instant observations become available, variational methods attempt to find a "best" model trajectory. The strong-constraint variational method ensures that over the entire time interval considered, the "best" model trajectory exactly obeys the model dynamics. Minimising a model-data misfit thereby becomes a constrained optimisa-

tion problem, with the constraint being the model dynamics. Weak-constraint variational methods, on the other hand, allow for some uncertainty in the model dynamics and do not require that the "best" trajectory is an exact solution of the model equations [*Losa et al.*, 2004]. In principle, the uncertainty in the model dynamics can also be accounted for by the strong-constraint method: by explicitly introducing error terms that can be turned on or off by adjustable parameters to be optimised during the optimisation process, a strong-constraint method can also account for (and quantify) model errors. For this reason, only strong-constraint variational methods will be discussed here.

The clue to solving constrained optimisation problems is to identify the so-called control parameters on which the solution, or trajectory, of the dynamical model depends. Such parameters may be initial conditions, boundary conditions (e.g., nutrient supply from outside the model domain), or internal model parameters (e.g., maximum growth rates, mortality rates). Together, they form a control parameter vector $\mathbf{p}$, and any particular choice of $\mathbf{p}$ will, for the model under consideration, uniquely determine the temporal evolution of the model trajectory. The total model data misfit over the considered time interval is then a function of $\mathbf{p}$ only. This function is often called cost function. Assuming a total of M observations d_j, $j = 1, \ldots, M$, and model counterparts m_j, one simple choice for the cost function is:

$$J(\mathbf{p}) = \sum_{j=1}^{M} [d_j - m_j(\mathbf{p})]^2 \tag{24.4}$$

Any prior information about the parameter values (e.g., physiological constraints, positiveness) or the model trajectory (e.g., possible deviations from a stationary seasonal cycle, smoothness) should enter the cost function in form of additional terms. Constructing appropriate terms should always be possible, and any information we cannot measure in this way is probably useless anyway.

In principle, things are easy now: finding the "best" model trajectory is equivalent to finding the parameter vector $\mathbf{p}^{opt}$ that minimises the cost function $J(\mathbf{p})$. At closer inspection, however, things are a little more complicated. In general, we will have different kinds of measurements, i.e., for different j the corresponding observations d_j (and model counterparts m_j) may have different units as well as different error statistics. This is usually dealt with by introducing a scale factor S_j and replacing $[d_j - m_j]$ by $[d_j - m_j]/S_j$. Various choices of S_j have been used so far, e.g., $S_j = \sigma(d_j)$, $S_j = d_j$, $S_j = m_j$, $S_j = d_{ave}$, $S_j = d_{max}$, and little emphasis is usually put on investigating the implications of the actual choise made [*Evans*, 2003].

Another issue is the functional form of the terms entering the cost function. The most common approach of using sums of squares in the cost function gives the same weight to positive and negative misfits, and minimising least squares corresponds to a maximum-likelihood estimate only as long as the error distributions are Gaussian. In most cases this will not be correct. For example, there is strictly zero probability that nutrient concentrations are negative. This could, for example, be reflected in the cost functions by terms that go to infinity as simulated nutrient concentrations approach zero (e.g., by a high negative power of the nutrient concentration). A further aspect to consider is a possible correlation of different observations in space and in time. Sometimes it is attempted to take these into account by weighting different observations by the number of measurements taken. More frequent observations are then assumed to be correlated and accordingly downweighted in the cost function, whereas rare observations get relatively more weight. If prior information on the parameter values is available, the respective cost-function terms have to be weighted against the model data misfit terms as well.

When all information is accounted for by appropriately weighted cost function terms, the resulting cost function defines the metrics that measures the qualitiy of any parameter set. Construction of the cost function will always involve some subjective elements regarding the functional forms or weights of the individual terms. This basically reflect that, as usual in life, different people have different views on what is "best". The optimisation results will always depend on the quality of the cost function which therefore should be crafted as carefully as possible.

5.3 Minimisation methods

In principle, one could just explore the cost function "landscape" in parameter space by explicitly evaluating $J(\mathbf{p})$ for a large number of different choices of the parameter vector $\mathbf{p}$. In practice, however, this will usually not be possible. Even for a simple NPZD-type ecosystem model with, say, 15 parameters, a very coarse sampling of only 10 possible values per parameter would require 10^{15} evaluations of the cost function. As soon as models have more than a handful parameters, more efficient minimisation methods are needed.

A large variety of such minimisation methods exist, most of which have been developed outside oceanography. They can be devided into methods that make use of the cost function's gradient, i.e., information about the downhill direction in the cost-function landscape, and into methods that do not use this information and therefore do not require

the often expensive computation of the gradient. In the following, a brief overview will be given of methods that have so far been applied to marine ecosystem models.

Gradient descent methods. A standard conjugate gradient method, that alters search directions in consecutive iterations, has been applied by *Fasham and Evans* [1995] to optimise a model at the site of the North Atlantic Bloom Experiment (47°N, 20°W). The cost function's gradient was approximated by varying the individual parameters by a finite amount and computing the corresponding difference quotient. A finite-difference approximation of the cost function gradient was also used by *Prunet et al.* [1996a,b]. It is generally and without detailed investigation assumed that the cost function is smooth enough so that the estimated gradient (times -1) points at least downward in the cost function topography and that errors in its exact direction and size will only slow, but not hinder, convergence of the descent algorithm. *Prunet et al.* [1996a,b] in addition made the probably more critical assumption of a locally parabolic shape of the cost function at each iteration of a gradient descent method. Their sensitivity analysis indicated that this method did not generally yield robust parameter estimates, and that posterior error estimates were too small compared to the results of their sensitivity experiments.

In order to improve the quality of the gradient computation, the adjoint method has received considerable attention. It was first used in the context of marine ecosystem models by *Lawson et al.* [1995]. The adjoint method computes the exact gradient of the cost function $J(\mathbf{p})$ by resorting to the method of Langrangian multipliers. This method has been widely used in statistical mechanics to derive the Euler-Lagrange equations. In essence, a Lagrange function L is defined as the cost function augmented by a additional terms that contain the model equations E_j multiplied by a corresponding (and a priori unknown!) Lagrangian parameter λ_j:

$$L(\mathbf{p}, \lambda, \mathbf{x}) = J(\mathbf{p}) + \sum_{j}^{j_{max}} \lambda_j E_j(\mathbf{x}) \quad . \tag{24.5}$$

At first sight, things now look much more complicated than for the minimisation of the cost function $J(\mathbf{p})$ only: The Lagrange function depends not only on the parameter vector $\mathbf{p}$ but also on a vector of Lagrangian multipliers λ and on the model state vector $\mathbf{x}$. However, because the function contains all the information we have, i.e., the cost function and the model dynamics, there are no further constraints to be considered, and the minimisation of the Lagrange function becomes an uncon-

strained problem. Accordingly, the minimum of $L(\mathbf{p}, \lambda, \mathbf{x})$ can "simply" be found by setting all its partial derivatives to zero. It turns out that the derivatives with respect to the components of the Lagrangian multipliers return the model dynamics, whereas the derivatives with respect to the components of the state vector will, after repeated application of the chain rule, return what is called the "adjoint model". The adjoint model can be understood as a model that runs the dynamics of the original "forward" model backward in time while being forced by the model-data misfits. A single backward run of the adjoint model, which uses the final state of a run of the forward model as initial condition, returns the full gradient of the cost function, $\nabla_{\mathbf{p}} J(\mathbf{p})$, at the position of the actual parameter vector $\mathbf{p}$. A gradient descent algorithm will then be needed to find a new parameter vector to start the next iteration of forward model run and adjoint model run.

The main advantage of the adjoint technique is its very efficient computation of a complete gradient in N-dimensional parameter space: Only a single forward and a single backward model integration are needed, whereas the other methods mentioned above require order N model integrations to compute an approximate version of the gradient. There is, of course, nothing like a free lunch: construction of the backward, or so-called adjoint model is a major effort. (Semi-)automatic compilers exist than can help turning the computer code of a forward model into its adjoint counterpart (e.g., described by *Marotzke et al.,* [1999]).

Applications of the adjoint method to parameter estimation for marine ecosystem models have shown some success [*Spitz et al.,* 1998, 2001; *Friedrichs,* 2002; *Gunson et al.,* 1999]. However, *Schartau et al.* [2001] reported that the likely existence of local minima in the cost function may require restarting the gradient search from a large variety of initial estimates of the parameter vector. This is a problem any gradient-descent minimisation method will have. So far, however, it has not been conclusively demonstrated that local minima of the cost function do indeed exist. They are very difficult to identify in N-dimensional parameter space. Visualisations of two-dimensional sections through the N-dimensional cost function topography often show local minima [*Athias et al.,* 2000; *Vallino,* 2000; *Schartau et al.,* 2001], but it is by no means clear whether these are also minima in the other $N-2$ directions.

Stochastic minimisation methods. In order to avoid the expensive and often cumbersome computation of the cost function's gradient and also to cope with possibly existing local minima without having to fully scan the complete parameter space in a "brute force" mode, search algorithms have been developed that contain stochastic elements. These

can explore large regions of the parameter space and reduce the chance of getting trapped in a local minimum early on. An example is the Markov Chain Monte Carlo method [*Harmon and Challenor*, 1997], that also addresses estimating the posterior error distribution of the optimal parameter values. Construction of other stochastic minimisation methods was guided by attempts to understand the emergence of apparently optimal structures in nature. One of these is the concept of simulated annealing that was used by *Matear* [1995] to optimise parameters of a suite of ecosystem models to observations at Station P in the subpolar North Pacific. The simulated annealing technique is analogous to the thermodynamics that describe cooling and cristallisation of liquids. It consists of an iterative random selection of state vectors within a slowly narrowing probability distribution around the "best" parameter vector of the previous iteration. The width of this probability distribution decreases with increasing iteration number as does the probability of new parameter vectors being accepted with a higher cost function value than the "best" one of the previous iteration. A finite probability of uphill steps in the cost-function landscape is required to escape local minima. The probability is formulated in terms of the Boltzman factor that describes energy fluctuations in statistical mechanics, where the probability of transitions to more energetic states increases with temperature. If the temperature is decreased slowly enough to avoid getting trapped in local minima, this can eventually lead to a liquid cristallising into a perfect lattice that represents the state with lowest possible potential energy. Replacing potential energy by the cost function value, this behaviour is mimicked by the minimisation algorithm with a "temperature" parameter in the Boltzman factor decreasing with increasing iteration number. For typical ecosystem-model applications, several tenthousand evaluations of the cost function are needed to obtain a robust result. Still, convergence to the global minimum cannot generally be proven.

Hurtt and Armstrong [1996] employed simulated annealing to minimise the model-data misfit for a new, implicitly size-structured ecosystem model at Bermuda and later [*Hurtt and Armstrong,* 1999] extended this approach to test simultaneous optimisation of a similar model at the distinct sites of the Bermuda Atlantic Time series Study (BATS, 32°N, 64°W) and Ocean Weather Ship India (59°N, 19°W). While they found that different ecological processes had to be considered at the two locations to achieve a reasonable fit, *Schartau and Oschlies* [2003a,b] reported that simultaneous optimisation of a NPZD model at three sites (BATS, OWS India, and the site of the North Atlantic Bloom Experiment at 47°N, 20°W) worked almost as well (or as badly!) as separate optimisations at the individual locations. Their optimisation method

of choice was a (micro-)genetic algorithm. Although this method also requires several tenthousand evaluations of the cost function, a genetic algorithm was found to be slightly more efficient than simulated annealing in an idealised model study by *Athias et al.* [2000]. The genetic algorithm basically looks after a population of parameter vectors $\mathbf{p}_i$ that are allowed to reproduce according to a fitness measured by their cost-function value $J(\mathbf{p}_i)$. Gene crossover in the reproduction step can be accounted for by exchanging various components of the parameter vector, and mutation can be included as a random perturbation of the individual parameter values in the reproduction step (Figure 3). Both

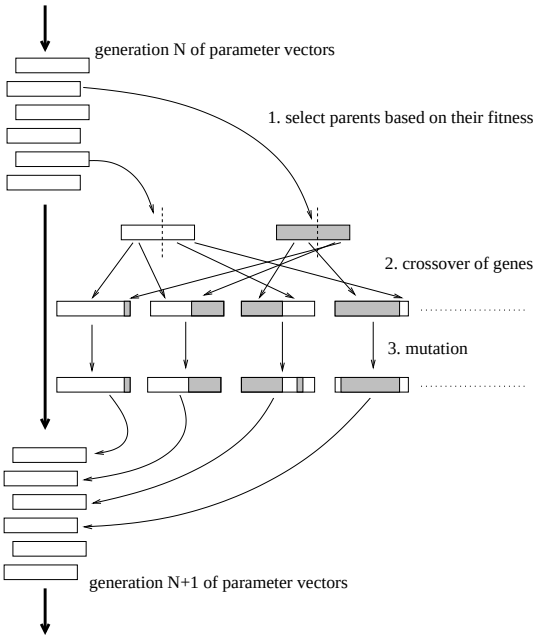

Figure 3. Schematic representation of the genetic algorithm. Out of an initial generation of parameter vectors, the fittest ones are selected according to their cost function value. In a second step, the selected parameter vectors can recombine (often the parameter vector is written as a binary string for this step), and in a third step there is low probability random mutation for individual elements of the parameter vector. For a micro-genetic algorithm, the fittest parameter vector of the parent generation is inherited to the next generation without any alterations. This elitism principle ensures that the lowest cost function value of each generation is at least as low as that one of the previous generation.

processes imply that even remote areas of the parameter space can be explored, whereas the gene-based reproduction concept ensures that information about good (i.e., "fit") parameter vectors (or their sub-sets)

is remembered by the algorithm. In many applications, the parameter "gene" is defined as the sequence of the components of the parameter vector written in binary notation, though a binary notation is not a necessary element of the method. As is the case with simulated annealing, convergence cannot generally be proven and will in practice depend on tunable parameters of the algorithm (mutation rate, population size).

6. Achievements and perspectives

The above assimilation and minimisation methods have so far been applied to ecosystem models that essentially all belong to the NPZD-type category. At present, there is no clear consensus on which method might be most efficient for this type of models. Sequential assimilation methods seem to have largest prospects in operational or near real-time applications that require good state estimates and that do not care that much about occasional jumps in the model trajectory. Variational methods, on the other hand, seem to be better suited for research issues that can take advantage of the dynamically self-consistent model trajectory. In principle, it is possible to obain a smooth model trajectory from sequential methods as well by applying a so-called smoother, which essentially consists of an integration backward in time. For linear systems it can be shown that sequential and variational algorithms can indeed produce identical results [*Bennett,* 1989].

A main advantage of stochastic optimisation techniques is the easy access to information about the posterior error of the parameter estimate. The large number of cost-function evaluations gives a reasonable picture of the cost function's sensitivity to changes in the individual parameters [e.g., *Schartau and Oschlies,* 2003a]. Particularly for strongly non-linear systems this may be more informative than local evaluations of the cost function's curvature via the Hessian matrix [*Fennel et al.,* 2001].

What is, to my knowledge, common to all of the assimilation studies performed so far is that at most 10 to 15 ecological model parameters could be constrained by the available observations. Although all these studies have so far employed relatively simple NPZD-type models, there were always a few parameters (or linear combinations of parameters) that could not be constrained. This indicates that even NPZD-type models have too many degrees of freedom and that models with fewer parameters should be able to reproduce the observations similarly well. Looking closer at how "well" the models can actually reproduce the data, one finds that even the optimised models fit the data very poorly. Usually, model-data misfits still amount to several standard deviations of the estimated prior errors. Such poor fits indicate that models need

more degrees of freedom to get closer to the observations. There are thus two contradictory statements about the required model complexity. This may to some extent result from errors in the physics used to drive the ecosystem models, but the problem is persistent even when physical observations are used to provide a physical environment as realistic as possible in a one-dimensional framework. Another possible explanation for the apparently contradictory statements about ecosystem model complexity is that the NPZD-type models employed so far may not have the right structures and hence are inherently inconsistent with the yet unknown ecological rules of real marine ecosystems. It is not clear whether more complex function-group type models would, in this data assimilative respect, perform any better. Although their order-of-magnitude larger number in adjustable parameters should allow for a much better fit to the data, such a fit would be of little value (and correspond to overfitting few data points by a high-order polynomial) for applying such models to other climate conditions unless all of the model parameters can be constrained by observations.

A possible strategy to clarify these issues is to undertake a systematic search for a model of minimum complexity that fits the available observations. Such an effort should start from a very simple model, perhaps similar to a nutrient-restoring one. Complexity should then be added only after careful analysis of the residual model-data misfits and some educated guess about the direction of complexity enhancement.

This approach for future model improvements should consider not only applications to the open ocean, where biogeochemical measurements are sparse and difficult to take, but also to controllable and manipulable field experiments, e.g., in artificial enclosures, and to physiological studies on cultures in the lab. A promising example is the application of ecosystem models and parameter optimisation methods to mesocosm experiments [*Vallino,* 2000]. Mesocosms are enclosed and generally gently mixed water volumes (typically several cubic meters in size), either in sea-water filled tanks on land, or in large plastic bags in the sea. They can be regarded as essentially homogeneous zero-dimensional systems that, in contrast to typical liter-sized incubation bottles, are large enough to keep boundary and enclosure effects small for a few generation times, i.e., days. By allowing for manipulation of environmental conditions and virtually unlimited access to observations, combining such experiments with modelling studies via data assimilative approaches may greatly help to better constrain our ecosystem models. On longer time scales, the same models will also have to be validated against open-ocean data collected within time-series programs and process studies. Physical models run in data assimilation mode may pro-

vide an optimal description of the physical environment for the marine ecosystem models to be run at these sites. Covering much longer time scales and more extreme climate states, even paleo information can help to constrain marine ecosystem models. By combining these very different time and space scales and the information provided by the different observational data sets with numerical models in an data assimilative approach, we will hopefully gain a better mechanistic understanding of marine ecosystem dynamics and their effects on biogeochemical cycles as well as their sensitivity to a changing climate.

Acknowledgments

I thank Eric Chassignet and Jacques Verron for allowing me to participate in this very enjoyable Summer School. This work is a contribution to the MERSEA Project. Partial support of the European Commission under contract SIP3-CT-2003-502885 is gratefully acknowledged.

References

Anderson, L. A., and J. L. Sarmiento, 1995: Global ocean phosphate and oxygen simulations. *Global Biogeochem. Cycles, 9*, 621–636.

Athias, V., P. Mazzega, and C. Jeandel, 2000: Selecting a global optimization method to estimate the oceanic particle cycling rate constants. *J. Mar. Res., 58*, 675–707.

Aumont, O., E. Maier-Reimer, S. Blain, and P. Monfray, An ecosystem model of the global ocean including Fe, Si, P co-limitations. *Global Biogeochem. Cycles, 17*(2), doi:10.1029/2001GB001745, 2003.

Bacastow, R., and E. Maier–Reimer, 1990: Ocean–circulation model of the carbon cycle. *Climate Dynamics, 4*, 95–125.

Bennett, A. F., 1989: The Kalman smoother for a linear quasi-geostrophic model of ocean circulation. *Dyn. Atmos. Oceans, 13*, 219–267.

Bertino, L., G. Evensen, and H. Wackernagel, 2003: Sequential data assimilation techniques in oceanography. *Int. Stat. Rev., 71,* 223–242.

Carmillet, V., J.-M. Brankart, P. Brasseur, H. Drange, G. Evensen, and J. Verron, 2001: A singular evolutive extended Kalman filter to assimilate ocean color data in a coupled physical-biochemical model of the North Atlantic ocean. *Ocean Modelling*, 3, 167–192.

Chai, F., R. T. Barber, and S. T. Lindley, 1996: Origin and maintenance of high nutrient condition in the equatorial Pacific. *Deep-Sea Res. II, 43*, 1031–1064.

Dring, M. J., and D. H. Jewson, 1982: What does ^{14}C uptake by phytoplankton really measure? A theoretical modelling approach. *Proc. R. Soc. Lond., B 214*, 351–368.

Ebenhöh, W., J. G. Baretta, and J. W. Baretta, 1997: The primary production module in the marine ecosystem model ERSEM II with emphasis on the light forcing. *J. Sea Res., 38,* 173–193.

Evans, G. T., 2003: Defining misfit between biogeochemical models and data sets. *J. Mar. Syst., 40–41,* 49–54.

Evensen, G., 1992: Using the extended Kalman filter with a multilayer quasi-geostrophic ocean model. *J. Geophys. Res., 97*(C11), 17,905–17,924.

Evensen, G., 2003: The Ensemble Kalman Filter: theoretical formulation and practical implementation. *Ocean Dynamics, 53*, 343–367.

Fasham, M. J. R., H. W. Ducklow, and S. M. McKelvie, 1990: A nitrogen–based model of plankton dynamics in the oceanic mixed layer. *J. Mar. Res., 48*, 591–639

Fasham, M. J. R., J. L. Sarmiento, R. D. Slater, H. W. Ducklow, and R. Williams, 1993: Ecosystem behavior at Bermuda Station "S" and Ocean Weather Station "India": a general circulation model and observational analysis. *Global Biogeochem. Cycles, 7*, 379–415.

Fasham, M. J. R., and G. T. Evans, 1995: The use of optimisation techniques to model marine ecosystem dynamics at the JGOFS station at 47N 20W. *Phil. Trans. Roy. Soc. Lond. B, 348*, 203–209.

Fennel, K., M. Losch, J. Schröter, and M. Wenzel, 2001: Testing a marine ecosystem model: sensitivity analysis and parameter optimization. *J. Mar. Syst., 28*, 45–63.

Friedrichs, M. A. M., 2002: Assimilation of JGOFS EqPac and SeaWiFS data into a marine ecosystem model of the central equatorial Pacific Ocean. *Deep-Sea Res. II, 49*, 289–319.

Gabric, A. J., R. Simo, R. A. Cropp, A. C. Hirst, and J. Dachs (2004), Modeling estimates of the global emission of dimethylsulfide under enhanced greenhouse conditions, *Global Biogeochem. Cycles, 18*, GB2014, doi:10.1029/2003GB002183.

Gunson, J., A. Oschlies, and V. Garçon, 1999: Sensitivity of ecosystem parameters to simulated satellite ocean colour data using a coupled physical-biological model of the North Atlantic. *J. Mar. Res., 57*, 613–639.

Harmon, R., and P. Challenor, 1997: A Markov chain Monte Carlo method for estimation and assimilation into models. Ecol. Model., 101, 41–59.

Hemmings, J. C. P., M. A. Srokocz, P. Challenor, and M. J. R. Fasham, 2003: Assimilating satellite ocean-colour observations into oceanic ecosystem models. *Phil. Trans. R. Soc. Lond., A 361,* 33–39.

Hurtt, G. C., and R. A. Armstrong, 1996: A pelagic ecosystem model calibrated with BATS data. *Deep-Sea Res. II, 43*, 653–683.

Hurtt, G. C., and R. A. Armstrong, 1999: A pelagic ecosystem model calibrated with BATS and OWSI data. *Deep-Sea Res. I, 46*, 27–61.

Kähler, P., and W. Koeve, 2001: Marine dissolved organic matter: can its C:N ratio explain carbon overconsumption? *Deep–Sea Res. I, 48*, 49–62.

Kalman, R. E., 1960: A new approach to linear filtering and prediction problems. *J. Basic. Eng., 82D,* 35–45.

Klausmeier, C. A., E. Litchman, T. Daufresne, and S. A. Levin, 2004: Optimal nitrogen-to-phosphorus stoichiometry of phytoplankton. *Nature*, 429, 171–174.

Körtzinger, A., W. Koeve, P. Kähler, and L. Mintrop, 2001: C:N ratios in the mixed layer during the productive season in the Northeast Atlantic Ocean. *Deep-Sea Res. I, 48*, 661–688.

Lawson, L. M., Y. H. Spitz, E. E. Hofmann, and R. B. Long, 1995: A data assimilation technique applied to a predator-prey model. *Bull. Math. Biology*, 57, 625–651.

Loukos, H., P. Monfray, L. Bopp, and P. Lehody, 2003: Potential changes in skipjack tuna (*Katsuwonus pelamis*) habitat from a global warming scenario: modelling approach and preliminary results. *Fisheries Oceanography, 12,* 474–482.

Losa, S. N., G. A. Kivman, and V. A. Ryabchenko, 2004: Weak constraint parameter estimation for a simple ocean ecosystem model: what can we learn about the model and data? *J. Mar. Syst., 45*, 1–20.

Marotzke, J., R. Giering, K. Q. Zhang, D. Stammer, C. Hill, and T. Lee, 1999: Construction of the adjoint MIT ocean general circulation model and application to Atlantic heat transport sensitivity. *J. Geophys. Res., 104*, 29,529–29,547.

Matear, R. J., 1995: Parameter optimization and analysis of ecosystem models using simulated annealing: a case study at Station P. *J. Mar. Res., 53*, 571–607.

McCreary, J. P., K. E. Kohler, R. R. Hood, and D. B. Olson, 1996: A four-component ecosystem of biological activity in the Arabian Sea. *Prog. Oceanog., 37*, 193–240.

Moore, J. K., S. C. Doney, J. A. Kleypas, D. M. Glover, and I. Y. Fung, 2002: An intermediate complexity marine ecosystem model for the global domain. *Deep–Sea Res. II, 49*, 403–462.

Najjar, R. G., J. L. Sarmiento, and J. R. Toggweiler, 1992: Downward transport and fate of organic matter in the ocean: simulations with a general circulation model. *Global Biogeochem. Cycles, 6*, 45–76.

Orr, J. C., 1999: Ocean Carbon-Cycle Model Intercomparison Project (OCMIP), *Research GAIM, 2*, 7.

Oschlies, A., 1999: On spurious interactions among a mixed layer model, convective adjustment, and isopycnal mixing in ocean circulation models. *Mon. Weather Rev., 127*, 1920–1927.

Oschlies, A., 2002: Can eddies make ocean deserts bloom? *Global Biogeochem. Cycles, 16*, 1106, doi:10.1029/2001GB001830.

Oschlies, A., 2004: Feedbacks of biotically induced radiative heating on upper-ocean heat budget, circulation, and biological production in a coupled ecosystem-circulation model. *J. Geophys. Res.,* 110, doi:10.1029/2004JC002430.

Oschlies, A., and V. Garçon, 1998: Eddy-induced enhancement of primary production in a model of the North Atlantic Ocean. *Nature, 394*, 266–269.

Oschlies, A., and V. Garçon, 1999: An eddy-permitting coupled physical-biological model of the North Atlantic. 1. Sensitivity to advection numerics and mixed layer physics. *Global Biogeochem. Cycles, 13*, 135–160.

Popova, E. E., M. J. R. Fasham, A. V. Osipov, and V. A. Ryabchenko, 1997: Chaotic behaviour of an ocean ecosystem model under seasonal external forcing. *J. Plank. Res., 19*, 1495–1515.

Popova, E. E., M. A. Srokosz, and D. A. Smeed, 2002: Real-time forecasting of biological and physical dynamics at the Island-Faeroes Front in June 2001. *Geophys. Res. Lett., 29*, 10.1029/2001GL013706.

Prunet, P., J. F. Minster, D. Ruiz–Pino, and I. Dadou, 1996: Assimilation of surface data in a one–dimensional physical–biogeochemical model of the surface ocean. 1. Method and preliminary results. *Global Biogeochem. Cycles, 10*, 111–138.

Prunet, P., J. F. Minster, V. Echevin, and I. Dadou, 1996: Assimilation of surface data in a one–dimensional physical–biogeochemical model of the surface ocean. 2. Adjusting a simple trophic model to chlorophyll, temperature, nitrate and P_{CO_2} data. *Global Biogeochem. Cycles, 10*, 139–158.

Redfield, A. C., 1934: On the proportions of organic derivatives in sea water and their relation to the composition of plankton. James Johnstone Memorial Volume, Liverpool, 176–192.

Redfield, A.C., B. H. Ketchum, F. A. Richards, 1963: The influence of organisms on the composition of sea water. In *The Sea*, Vol. 2, M. N. Hill (Ed.), Interscience, New York, 26–77.

Sarmiento, J. L., R. D. Slater, M. J. R. Fasham, H. W. Ducklow, J. R. Toggweiler, and G. T. Evans, 1993: A seasonal three–dimensional ecosystem model of nitrogen

cycling in the North Atlantic euphotic zone. *Global Biogeochem. Cycles, 7*, 417–450.

Sarmiento, J. L., R. Murnane, and C. Le Quéré, 1995: Air-Sea CO_2 transfer and the carbon budget of the North Atlantic. *Philos. Trans. R. Soc. London Ser. B, 348*, 211–219.

Schartau, M., A. Oschlies, and J. Willebrand, 2001: Parameter estimates of a zero-dimensional ecosystem model applying the adjoint method. *Deep-Sea Res. II, 48*, 1769-1800.

Schartau, M., and A. Oschlies, 2003: Simulataneous data-based optimization of a 1D-ecosystem model at three locations in the North Atlantic: Part I–Method and parameter estimates. *J. Mar. Res., 61*, 765–793.

Schartau, M., and A. Oschlies, 2003: Simulataneous data based optimization of a 1D-ecosystem model at three locations in the North Atlantic: Part II–Standing stocks and nitrogen fluxes. *J. Mar. Res., 61*, 795–821.

Schofield, O., J. Grzymski, W. P. Bissett, G. J, Kirkpatrick, D. F. Millie, M. Moline, and C. S. Roesler, 1999: Optical monitoring and forecasting systems for harmful algal blooms: possibility of pipe dream? *J. Phycol., 35*, 1477–1496.

Six, K. D., and E. Maier-Reimer, 1996: Effects of plankton dynamics on seasonal carbon fluxes in an ocean general circulation model. *Global Biogeochem. Cycles, 10*, 559–583.

Smetacek, V., 2001: A watery arms race. *Nature, 411*, 745.

Spitz, Y. H., J. R. Moisan, M. R. Abbott, and J. G. Richman, 1998: Data assimilation and a pelagic ecosystem model: parameterization using time series observations. *J. Mar. Syst., 16*, 51–68.

Spitz, Y. H., J. R. Moisan, and M. R. Abbott, 2001: Configuring an ecosystem model using data from the Bermuda Atlantic Time Series (BATS). *Deep–Sea Res. II, 48*, 1733–1768.

Sterner, R. W., and J. J. Elser, 2002: *Ecological Stoichiometry: the biology of elements from molecules to the biosphere*, Princeton Univ. Press, Princeton, 439pp.

Vallino, J. J., 2000: Improving marine ecosystem models: Use of data assimilation and mesocosm experiments. *J. Mar. Res., 58*, 117–164.

Volk, T., and M. I. Hoffert, 1985: Ocean carbon pumps: Analysis of relative strengths and efficiencies in ocean-driven atmospheric CO_2 changes, in *The Carbon Cycle and Atmospheric CO_2: Natural Variations, Archean to Present*, edited by E. Sundquist and W. Broecker, pp. 99–110, AGU Geophysical Monograph, 32, Washington D.C.

Williams, P. L. le B., 1995: Evidence for the seasonal accumulation of carbon-rich dissolved organic material, its scale in comparison with changes in particulate material and the consequential effect on net C/N assimilation rates. *Mar. Chem., 51*, 17–29.